Presented at IMAC 18, Feb 7-10, 2000, San Antonio, Texas.

ISSUES FOR THE APPLICATION OF STATISTICAL MODELS IN DAMAGE DETECTION

Matthew T. Bement[1], Charles R. Farrar[2]

[1]Department of Mechanical Engineering, MS 3123, Texas A&M University, College Station, TX, 77843
[2]Engineering Analysis Group, MS P946, Los Alamos National Laboratory, Los Alamos, NM, 87545

ABSTRACT

Many aerospace, civil, and mechanical systems continue to be used despite aging and the associated potential for damage accumulation. Therefore, the ability to monitor the structural health of these systems is becoming increasingly important. A wide variety of highly effective local non-destructive evaluation tools are available. However, damage identification based upon changes in vibration characteristics is one of the few methods that monitor changes in the structure on a global basis. The process of vibration-based damage detection will be described as a problem in statistical pattern recognition. This process is composed of four portions: 1.) Operational Evaluation, 2.)Data acquisition and cleansing; 3.) Feature selection and data compression, and 4.) Statistical model development. Current studies regarding supervised learning methods for statistical model development are discussed and emphasized with the application of this technology to a laboratory test structure. Specifically, a comparison is made between a linear discriminant classifier and a general Bayesian classifier for the purpose of determining the existence of damage.

1. INTRODUCTION

In very general terms damage can be defined as changes introduced into a system that adversely affect the current or future performance of that system. Implicit in this definition is the concept that damage is not meaningful without a comparison between two different states of the system, one of which is assumed to represent the initial, and often undamaged, state. This discussion is focused on the study of damage identification in structural and mechanical systems. Therefore, the definition of damage will be limited to changes to the material and/or geometric properties of these systems, including changes to the boundary conditions and system connectivity, which adversely effect the current or future system performance.

The interest in the ability to monitor a structure and detect damage at the earliest possible stage is pervasive throughout the civil, mechanical and aerospace engineering communities. Current damage-detection methods are either visual or localized experimental methods such as acoustic or ultrasonic methods, magnetic field methods, radiograph, eddy-current methods and thermal field methods (Doherty, 1987). All of these experimental techniques require that the vicinity of the damage be known *a priori* and that the portion of the structure being inspected is readily accessible. Subjected to these limitations, these experimental methods can detect damage on or near the surface of the structure. The need for quantitative global damage detection methods that can be applied to complex structures has motivated research of methods that examine changes in the vibration properties of the structure.

The basic premise of vibration-based damage detection is that the damage will significantly alter the stiffness, mass or energy dissipation properties of a system, which, in turn, will alter the measured dynamic response of that system. Although the basis for vibration-based damage detection appears intuitive, its actual application poses many significant technical challenges. The most fundamental challenge is the fact that damage is typically a local phenomenon and may not significantly influence the lower-frequency global response of structures that is typically measured during vibration tests. This challenge is supplemented by many practical issues associated with making accurate and repeatable vibration measurements at a limited number of locations on structures often operating in adverse environments.

Recent research has begun to recognize that the vibration-based damage detection problem is fundamentally one of statistical pattern recognition and this paradigm is described in detail. In particular, the study reported herein provides a comparison of two pattern classification methods.

2. HISTORICAL PERSPECTIVE

The development of vibration-based damage detection technology has been closely coupled with the evolution, miniaturization and cost reductions of Fast Fourier Transform (FFT) analyzer hardware and computing hardware. To date, the most successful application of vibration-based damage detection technology has been for monitoring rotating machinery. The detection process is based on pattern recognition applied to time histories or spectra generally measured on the housing of the machinery during normal operating conditions.

The aerospace community began to study the use of vibration-based damage detection during the late 1970's and early 1980's in conjunction with the development of the space shuttle. The Shuttle Modal Inspection System (SMIS) was developed to identify fatigue damage in components such as control surfaces, fuselage panels and lifting surfaces. This system has been successful in locating damaged components covered by the thermal protection system, and all orbiter vehicles have been periodically subjected to SMIS testing since 1987.

The civil and petroleum engineering communities have studied vibration based damage assessment for large scale structures such as bridge structures and offshore drilling platforms. Difficulties associated with the changing, yet undamaged structural properties of offshore drilling platforms ended the petroleum industry's interest in vibration based damage assessment for drilling platforms in the late 80's. However, regulatory requirements in Asian countries, which mandate the companies that construct bridges periodically certify their structural health, are driving current research and development of vibration-based bridge monitoring systems.

In summary, the review of the technical literature presented by (Doebling et al. 1996) [1] shows an increasing number of research studies related to vibration-based damage detection. These studies identify many technical challenges to the adaptation of vibration-based damage detection that are common to all applications of this technology. These challenges include better utilizing the nonlinear response characteristics of the damaged system, development of methods to optimally define the number and location of the sensors, identifying the features sensitive to small damage levels, the ability to discriminate changes in features cause by damage from those caused by changing environmental and/or test conditions, the development of statistical methods to discriminate features from undamaged and damaged structures, and performing comparative studies of different damage detection methods applied to common data sets. These topics are currently the focus of various research efforts by many industries including defense, automotive, and semiconductor manufacturing where multi-disciplinary approaches are being used to advance the current capabilities of vibration-based damage detection.

3. VIBRATION-BASED DAMAGE DETECTION AND STRUCTURAL HEALTH MONITORING

The process of implementing a damage detection strategy is referred to as *structural health monitoring*. This process involves the observation of a structure over a period of time using periodically spaced measurements, the extraction of features from these measurements, and the analysis of these features to determine the current state of health of the system. The output of this process is periodically updated information regarding the ability of the structure to continue to perform its desired function in light of the inevitable aging and degradation resulting from the operational environments.

3. 1. Operational Evaluation

Operational evaluation answers four questions in the implementation of a structural health monitoring system:

1. How is damage defined for the system being studied?
2. What are the economic and/or life safety justification for performing the health monitoring activity?
3. What are the conditions, both operational and environmental, under which the system to be monitored functions?
4. What are the limitations on acquiring data in the operational environment?

Operational evaluation begins to set the limitations on what will be monitored, why will it be monitored, and how the monitoring will be accomplished. This evaluation starts to tailor the damage detection process to features that are unique to the system being monitored and tries to take advantage of unique features of the postulated damage that is to be detected.

3. 2. Data Acquisition and Cleansing

The data acquisition portion of the structural health monitoring process involves selecting the types of sensors to be used, the location where the sensors should be placed, the number of sensors to be used, and the data acquisition/storage/transmittal hardware. This process will be application specific. Economic considerations will play a major role in making these decisions. Another consideration is how often the data should be collected. In some cases it may be adequate to collect data immediately before and at periodic intervals after a severe event. However, if fatigue crack growth is the failure mode of concern, it may be necessary to collect data almost continuously at relatively short time intervals.

Because data can be measured under varying conditions, the ability to normalize the data becomes very important to the damage detection process. One of the most common procedures is to normalize the measured responses by the measured inputs. When environmental or operating condition variability is an issue, the need can arise to normalize the data in some temporal fashion to facilitate the comparison of data measured at similar times of an environmental or operational cycle. Sources of variability in the data acquisition process and with the system being monitored need to be identified and minimized to the extent possible. In general, all sources of variability can not be eliminated. Therefore, it is necessary to make the appropriate measurements such that these sources can be statistically quantified.

Data cleansing is the process of selectively choosing data to accept for, or reject from, the feature selection process. The data cleansing process is usually based on knowledge gained by individuals directly involved with the data acquisition. One of the most common forms of data cleansing is to apply various filters to the data. Finally, it should be noted that the data acquisition and cleansing portion of a structural health-monitoring process should not be static. Insight gained from the feature selection process and the statistical modeling process will provide information that can improve the data acquisition process.

3. 3. Feature Selection

The area of the structural damage detection process that receives the most attention in the technical literature is the identification of data features that allow one to distinguish between the undamaged and damaged structure. Inherent in this feature selection process is the condensation of the data. The operational implementation and diagnostic measurement technologies needed to perform structural health monitoring typically produce a large amount of data. A condensation of the data is advantageous and necessary particularly if comparisons of many data sets over the lifetime of the structure are envisioned. Also, because data may be acquired from a structure over an extended period of time and in an operational environment, robust data reduction techniques must retain sensitivity of the chosen features to the structural changes of interest in the presence of environmental noise.

The best features for damage detection are typically application specific. Numerous features are often identified for a structure and assembled into a feature vector. In general, it is desirable to develop feature vectors that are of low dimension. It is also desirable to obtain many samples of the feature vectors. There are no restrictions on the types or combinations of data contained in the feature vector. As an example, a feature vector may contain the first three resonant frequencies of the system, a time when the measurements were made, and a temperature reading from the system.

A variety of methods are employed to identify features for damage detection. Past experience with measured data from a system, particularly if damaging events have been previously observed for that system, is often the basis for feature selection. Numerical simulation of the damaged system's response to simulated inputs is another means of identifying features for damage detection. The application of engineered flaws, similar to ones expected in actual operating conditions, to specimens can identify parameters that are sensitive to the expected damage. Damage accumulation testing, during which significant structural components of the system under study are subjected to a realistic accumulation of damage, can also be used to identify appropriate features. Fitting linear or nonlinear, physical-based or non-physical-based models of the structural response to measured data can also help identify damage-sensitive features. A detailed summary of features that have been used for vibration-base damage detection can be found in (Doebling, et al., 1996) [1].

3. 4. Statistical Model Development

The portion of the structural health monitoring process that has received the least attention in the technical literature is the development of statistical models to enhance the damage detection. Almost none of the hundreds of studies summarized in (Doebling, et al, 1996) [1] make use of any statistical methods to assess if the changes in the selected features used to identify damaged are statistically significant. Statistical model development is concerned with the implementation of the algorithms that operate on the extracted features and unambiguously determine the damage state of the structure. The algorithms used in statistical model development usually fall into three categories and will depend on the availability of data from both an undamaged and damaged structure. The first category is group classification, that is, placement of the features into respective "undamaged" or "damaged" categories. Analysis of outliers is the second type of algorithm. When data from a damaged structure are not available for comparison, do the observed features indicate a significant change from the previously observed features that can not be explained by extrapolation of the feature distribution? The third category is regression analysis. This analysis refers to the process of correlating data features with particular types, locations or extents of damage. All three algorithm categories analyze statistical distributions of the measured or derived features to enhance the damage detection process.

The statistical models are used to answer the following questions regarding the damage state of the structure: 1. Is there damage in the structure (existence)?; 2. Where is the damage in the structure (location)?; and 3. How severe is the damage (extent)? Answers to these questions in the order presented represents increasing knowledge of the damage state. Experimental structural dynamics techniques can be used to address the first two questions. Analytical models are usually needed to answer the third question unless examples of data are available from the system (or a similar system) when it exhibits varying level of the damage. Statistical models can also be used to determine the type of damage that is present. To identify damage type, data from damaged structures must be available for correlation with the measured features.

Finally, an important part of the statistical model development process is the testing of these models on actual data to establish the sensitivity of the selected features to damage and to study the possibility of false indications of damage. False indications of damage fall into two categories: 1.) False-positive damage indication (indication of damage when none is present), and 2). False-negative damage indications (no indication of damage when damage is present). Although the second category is usually very detrimental to the damage detection and can have serious life-safety implications, false-positive readings can also erode confidence in the damage detection process.

This paper will now summarize the application of methods from statistical pattern recognition and machine learning to a vibration-based damage detection problem. A damage detection experiment performed on an 8-DOF system will be described in terms of the statistical-pattern-recognition damage-detection paradigm just summarized.

4. SOME APPROACHES TO SUPERVISED LEARNING

Consider the process of classifying data into one of two classes, denoted as A and B. In the case of supervised learning, it is assumed that several examples of data belonging to each class are available. The goal is then to use the examples to determine which class a new piece of data should be assigned to. While there are many different methods for accomplishing this task, we limit our consideration to two methods that are statistical in nature. The first utilizes a linear discriminant method known as Fisher's discriminant [2] to determine the probability that a new data point belongs to a given class. The second method, Bayesian classification [3], also aims to predict the probability that a new data point belongs to a given class, but it is somewhat more general. Both methods have their advantages and disadvantages, as will be seen.

4.1 Fisher's Discriminant

In effect, Fisher's discriminant projects the two classes onto a line through the origin in the n-dimensional feature space such that the separation between the classes is maximized, while accounting for the in-class and between class scatter in the data sets. To be more precise, the vector *w* which lies along the desired line is determined such that

$$F(w) = \frac{(m_A - m_B)^2}{s_A^2 + s_B^2} \tag{1}$$

is maximized, where $y_k = w^T x_k$, $m_k = mean(y_k)$, the within class variance s_k^2 is given by $s_k^2 = \sum_{n \in k} (y_n - m_k)^2$, and $k \in \{A, B\}$. Once *w* is determined, and assuming that the prior distribution for *A*, *f(A)* (a measure of how frequently *A* is to be expected) is discrete, the probability that the transformed new data, $y_{new} = w^T x_{new}$ came from class *A* is given by application of Bayes theorem

$$f(A|y_{new}) = 1 - f(B|y_{new}) = \frac{f(y_{new}|A) f(A)}{f(y_{new}|A) f(A) + f(y_{new}|B) f(B)} \tag{2}$$

where $f(\cdot)$ is the appropriate probability density function.

4.2 Bayesian classification

Bayesian classification is more general than Fisher's discriminant. In fact Bayesian classification is used in the method described above, but only on the transformed data. In the more general case, the probability that the new data x_{new} belongs to class *A* is

$$f(A|x_{new}) = 1 - f(B|x_{new}) = \frac{f(x_{new}|A) f(A)}{f(x_{new}|A) f(A) + f(x_{new}|B) f(B)} \tag{3}$$

This is obviously very similar to the equation presented in the above section. The key difference is that the probability density functions are now multivariate, rather than univariate, since x_{new} is a vector. While this does present added computational difficulty, it should be noted that while Fisher's discriminant can be used only in the two class problem, the general Bayesian classification can be used in an (up to) infinite class problem.

4.3 Comparison

Note that in both the general Bayesian classification as well as in Bayesian classification after application of Fisher's discriminant, one must either assume a distribution for each class or determine one empirically in order to evaluate the $f(x_{new}|\cdot), f(y_{new}|\cdot)$ terms. In the case where Fisher's discriminant is used, y_{new} is merely a linear combination of the individual features. Thus, we have some justification for assuming a normal distribution for the projection of each class by appealing to the central limit theorem. However, in multidimensional Bayesian classification, one should either make sure that the assumed distribution for each class is justified, or one should determine the distribution empirically. Despite these difficulties, there are cases where the complications of using the general Bayesian classification are justified.

Figure 1 demonstrates a bivariate two-class problem. It illustrates a potential pitfall in using Fisher's discriminant. That is, if several new data points are to be classified, and they happen to lie on a line orthogonal to the line that provides maximum discrimination, the results of the classification may not be very helpful. This admittedly pathological example does not pose a problem in the general Bayesian classification. For this example, both class A and class B can be characterized by bivariate Gaussian distributions.

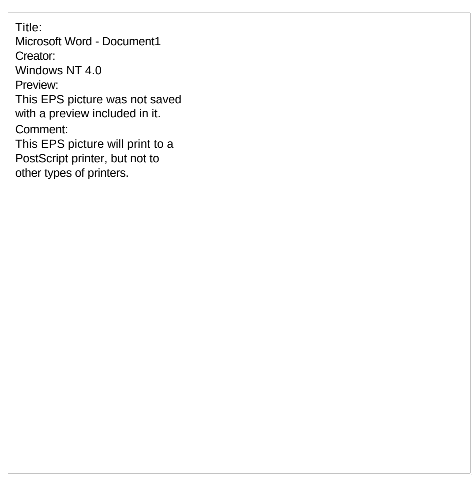

Fig. 1 Hypothetical bivariate two class problem

The Table 1 lists the probability that the new data points came from class B, using both Fisher's discriminant and a general Bayesian classification.

Table 1 Classification Probabilities for New Data in Figure 1		
Data Point	With Fisher	Without Fisher
1	0.5401	0.2925
2	0.5401	0.5650
3	0.5401	0.9165
4	0.5401	0.9960

As expected, with the application of Fisher's discriminant, all the points have the same probability of coming from class B. Thus, it would be difficult to determine with any great confidence which class any of the four points came from. In the general multidimensional classification, on the other hand, one would feel more comfortable assigning the new points to classes for all points but point number 2.

5. APPLICATION TO 8-DOF SYSTEM

In an effort to judge the performance of the statistical classification methods described above in a real world system, an experiment was performed to attempt to classify a system as being damaged or undamaged, based on its vibration response.

5.1 Experiment description

The system under consideration consisted of eight masses in series, connected with springs. A schematic of the system in a typical configuration is shown in Fig. 2 below.

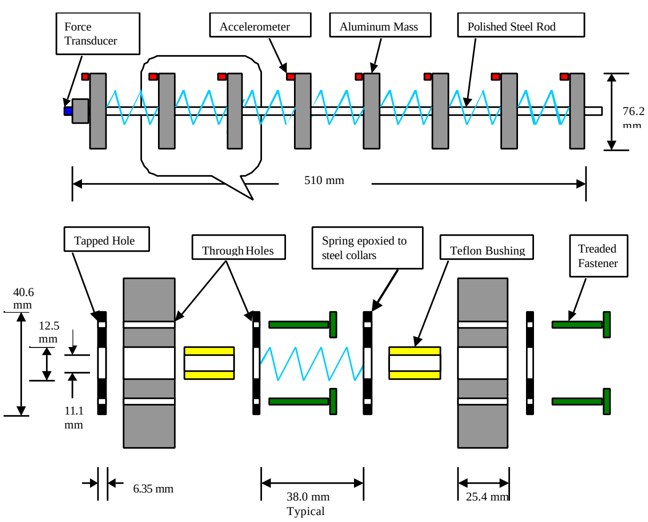

Fig. 2 Eight DOF test system

For the experiment, a removable "bumper" was installed in between the fifth and sixth mass, so as to limit the compression of the spring. When the bumper was removed, the system was considered undamaged. When the bumper was present, the system was considered damaged. The system was excited by a random signal produced by a shaker that was attached to the first mass. Accelerometers were attached to the first and sixth masses, and their outputs were recorded for each of the five, eight second trials that were performed for both the damaged and undamaged cases. A force transducer between the shaker and the first mass recorded the input force supplied by the shaker.

5.2 Feature selection

Each of the eight second trials was divided into eight, one second windows containing 512 points. While perhaps not strictly true, each of these windows was viewed as a statistically independent sample. Thus, the experiment yielded 40 examples of an undamaged response and 40 examples of a damaged response. For both the undamaged and damaged cases, 32 of the responses were used for training purposes, with the remaining eight being saved for validation. All eight of the validation responses came from the same trial. The trial chosen for validation was varied, as will be discussed below.

Typical undamaged and damaged time responses are shown in the Figs. 3 and 4 below.

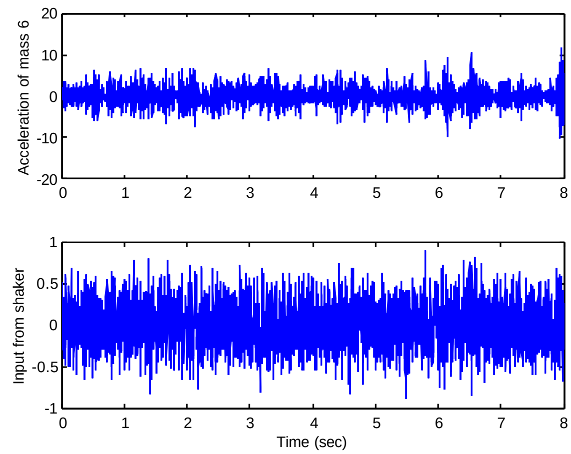

Fig. 3. Representative undamaged response

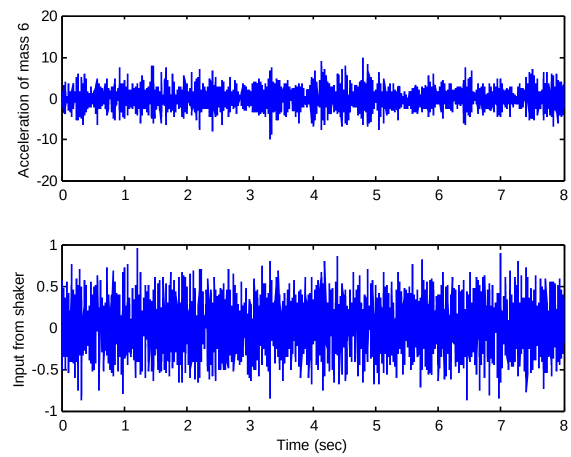

Fig. 4. Representative damaged response

Next, it was necessary to identify a few important features in the data, to allow classification. Since the number of data points (32) in each training set is relatively small, we would like to have a relatively small number of features. Thus, we considered auto-regressive (AR) and auto-regressive exogeneous (ARX) coefficients [4], because of their ability to characterize a system with a relatively small number of parameters. For the actual calculation of the coefficients, a 512 point hamming window was applied to all of the 512 point input and output samples. This was done to minimize the "end effects" in calculating the various coefficients.

After using the Matlab System Identification toolbox [5], an eighth order AR model and a seventh order ARX model (na = 4, nb = 3, nk=7) were decided on. Thus, each set of coefficients represents a point in the feature space. Thus, both the undamaged and damaged classes contained 32 points in the 8 (AR) or 7 (ARX) dimensional feature spaces.

For convenience, Gaussian distributions were assumed for both methods of classification, with the means and standard deviations of the distributions having been determined from the respective training sets.

5.3 Results of the classification

The following table shows the results of the classification for the AR model. The trial used for validation was varied to obtain some estimate of how sensitive the results were to the training sets. The column labeled “Damaged” gives the probability that the damaged validation trial belongs to the undamaged class. Likewise, the column labeled “Undamaged” gives the probability that the undamaged validation trial belongs to the undamaged class. Ideally, we would like to see all entries be “1”.

Trial # used for validation	Damaged (w/ Fisher)	Damaged (w/o Fisher)	Undamaged (w/ Fisher)	Undamaged (w/o Fisher)
1	1	0.9997	1	1
2	1	1	1	1
3	1	1	1	1
4	1	1	1	1
5	1	1	1	1

The following table shows the same information as the above table, but for the ARX model.

Trial # used for validation	Damaged (w/ Fisher)	Damaged (w/o Fisher)	Undamaged (w/ Fisher)	Undamaged (w/o Fisher)
1	1	1	1	0
2	1	1	1	1
3	1	1	1	1
4	1	1	1	1
5	1	1	1	1

Clearly, both methods of classification yield fairly good results that are not very sensitive to the trials that were chosen as training sets. There was one false positive when the ARX model was used in the general Bayesian classification. This result is most likely due to the assumption of Gaussian distributions, which was made only for convenience. For a more accurate picture of the capabilities of the general Bayesian classification in this particular problem, one should determine the distribution empirically. The consequences of the central limit theorem apparently prevented a similar problem from occurring for the classification after application of Fisher’s discriminant.

6. SUMMARY AND IMPORTANT ISSUES

For the experimental example presented in this paper, application of the Fisher discriminant resulted in the best classification results, in that no false or even ambiguous classifications occurred. However, even by making the assumption of Gaussian distributions (which had no empirical basis), the generalized Bayesian classification did a reasonable job of classification. Theoretically, the generalized classification should give the best results, provided that accurate distributions for the training sets can be obtained. Thus, before drawing any definite conclusions about which is better, it would be appropriate to estimate the multivariate distributions empirically. Also while applying Fisher’s discriminant is computationally appealing, it can only be used for two class problems, whereas the general multidimensional classification allows any number of classes to be considered.

Another very important issue that strongly affects the results of the classification is the choice of features. For this experiment we chose AR and ARX coefficients, which essentially fit a linear model to the system. However, when dealing with damage scenarios that are fundamentally nonlinear, other features might be necessary, especially if one hopes to ascertain the location and/or extent of the damage.

Finally, there still exists an important question when doing any kind of classification of this type. Are there really two (or more) distinct classes? This question could be answered by application of any number of methods in cluster analysis or self-organized learning. Answering this question is important in determining how much confidence one would have in the results of the classification.

ACKNOWLEDGEMENT

Portion of the funding for this work has come from a cooperative research and development agreement with Kinemetrics Corporation, Pasadena California. The civilian application of this CRADA is aimed at developing structural health monitoring systems for civil engineering infrastructure. A portion of the funding for this research was provided by the Department of Energy’s Enhanced Surveillance Program.

REFERENCES

[1] Doebling, S. W., Farrar, C. R., Prime , M. B., and Shevitz, D. W., "A Review of Damage Identification Methods that Examine Changes in Dynamic Properties," *Shock and Vibration Digest* **30** (2), pp. 91-105, 1998.

[2] Bishop, C. M., *Neural Networks for Pattern Recognition*, Oxford University Press, Oxford, UK, 1995.

[3] Freund, J. E. *Mathematical Statistics*, 5^{th} Edition, Prentice Hall, 1992.

[4] Juang, J-N, *Applied System Identification*, Prentice Hall, 1994.

[5] The Math Works, Inc., *System Identification Toolbox*, Natick, MA., 1995.